

Strategic Supercomputing Comes of Age

As ASC supercomputers

settle into the terascale


regime and prepare for

petascale, simulations

reveal new physics

insights to

scientists.



As the National Nuclear Security Administration's (NNSA's) Advanced Simulation and Computing (ASC) Program prepares to move into its second decade, the users of ASC's enormous computers also prepare to enter a new phase. Since its beginning in 1995, the ASC Program (originally the Accelerated Strategic Computing Initiative, or ASCI) has been driven by the need to analyze and predict the safety, reliability, and performance of the nation's nuclear weapons and certify their functionality—all in the absence of nuclear weapons testing. To that end, Lawrence Livermore, Los Alamos, and Sandia national laboratories have worked with computer industry leaders such as IBM, Intel, SGI, and Hewlett Packard to bring the most advanced and powerful machines to reality.

But hardware is only part of the story. The ASC Program also required the development of a computing infrastructure and scalable, high-fidelity, three-dimensional simulation codes to address issues related to stockpile stewardship. Most important, the laboratories had to provide proof of principle that users could someday have confidence in the results of the simulations when compared with data from legacy codes, past nuclear tests, and nonnuclear science experiments.

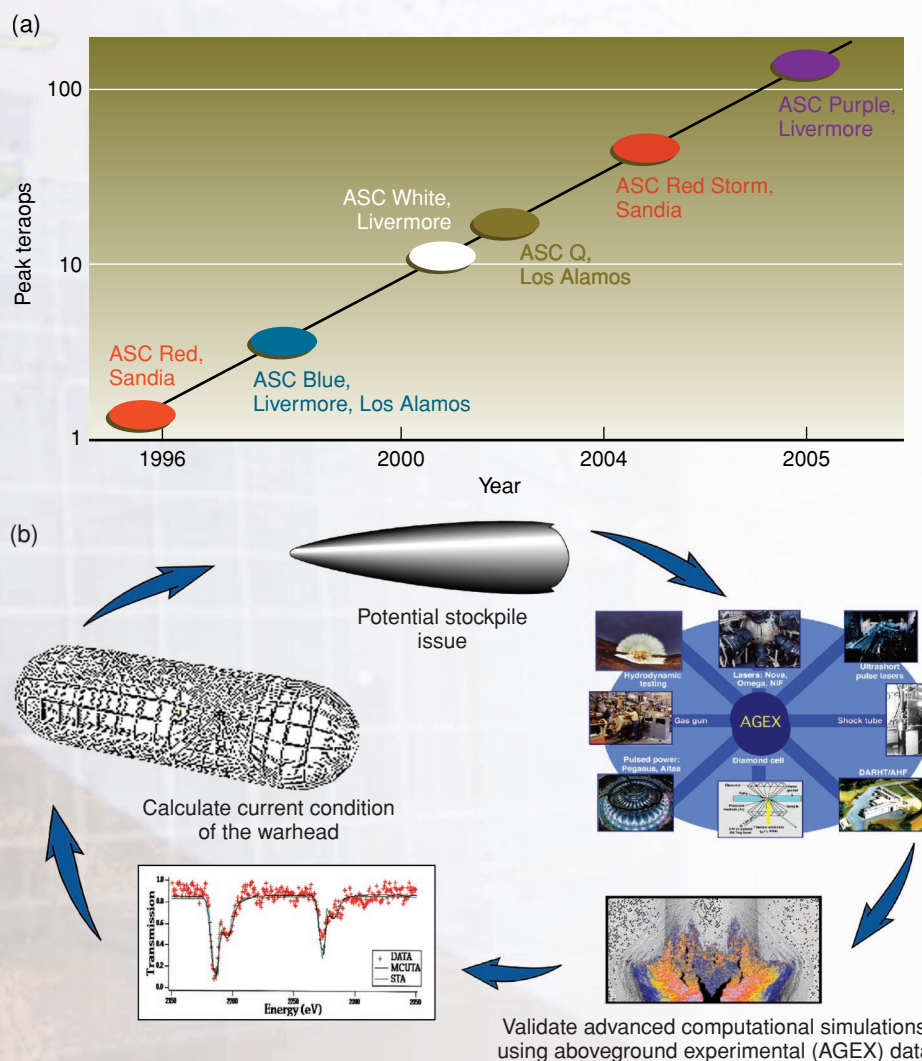
Efforts are now successfully moving beyond that proof-of-principle phase, notes Randy Christensen, who leads program planning in the Defense and Nuclear Technologies (DNT) Directorate and is one of the founding members of the tri-laboratory ASC Program. Christensen says, "With the codes, machines, and all the attendant infrastructure in order, we can now advance to the next phase and focus on improving the physics models in our codes to enhance our understanding of weapons behavior." Livermore's 12.3-teraops (trillion operations per second) ASC White machine and Los Alamos's 20-teraops ASC Q machine are in place, and the next systems in line are Sandia's 40-teraops ASC Red Storm and Livermore's 100-teraops ASC Purple. "In anticipation of ASC Purple in 2005, we are shifting our emphasis from

developing parallel-architecture machines and codes to improved weapons science and increased physics understanding of nuclear weapons," adds Christensen. "We are taking the next major step in the road we mapped out at the start of the program."

A Long and Winding Road

"Ten years ago, we were focused on creating a new capability, and the program was viewed more as an experiment or an

initiative," says Mike McCoy, acting leader for DNT's ASC Program. "Many skeptics feared that the three-dimensional codes we were crafting, and the new machines we needed to run them on, would fail to be of use to the weapons program." These skeptics had three areas of concern: First, would the new three-dimensional codes be useful? That is, would the code developers, working with other scientists, be able to develop new applications with the physics, dimensionality,



(a) The first seven years of the Accelerated Strategic Computing Initiative focused on creating a new capability for the Stockpile Stewardship Program, including the demonstration of three-dimensional applications. ASC Purple, slated to arrive in 2005, is the "entry point" for conducting high-fidelity, full weapons system simulations. (b) Over the next 10 years, the focus will be on using those capabilities as an integral part of the Stockpile Stewardship Program and on improving predictive capabilities.

resolution, and computational speed needed to take the next step in predictivity? Second, would the computers be reliable and work sufficiently well to grind through the incredibly complex and detailed calculations required in a world without underground nuclear testing? Third, would the supporting software infrastructure, or simulation environment, be able to handle the end-to-end computational and assessment processes? For that first decade, the program's primary focus was on designing codes and running prototype problems to address these concerns.

"Sophisticated weapon simulation codes existed before the ASC and Stockpile Stewardship programs," says Christensen. "However, because of the limited computer power available, those codes were never expected to simulate all the fine points of an exploding nuclear weapon. When the results of these simulations didn't match the results of the underground tests, numerical 'knobs' were tweaked to make the simulation results better match the experiments. When underground nuclear testing was halted in 1992, we could no

longer rely so heavily on tweaking those knobs."

At the time that underground testing ceased and NNSA's Stockpile Stewardship Program was born, Livermore weapons scientists were depending on the (then) enormous machines developed by Seymour Cray. Cray designed several of the world's fastest vector-architecture supercomputers and introduced closely coupled processors. "We had reached the limits on those types of systems," says McCoy, who is also a deputy associate

Not Just Computers and Codes: Making It All Work

Designers and physicists in the tri-laboratory (Livermore, Los Alamos, and Sandia national laboratories) Advanced Simulation and Computing (ASC) Program are now using codes and supercomputers to delve into regimes of physics heretofore impossible to reach. What made these amazing tools possible were the efforts of the computer scientists, mathematicians, and computational physicists who brought the machines and the codes to the point of deployment.

It wasn't easy. Throughout the era of testing nuclear weapons, approximations were a given for the computations. When calculations produced unusual results, scientists assumed that lack of resolution or faithful replication of geometry or faithful physics models or some combination were the culprits. "It was assumed that, no matter how big the machines were at that time, this inaccuracy would remain a given," says Mike McCoy, deputy associate director of Livermore's Computation Directorate. "But this concern was greatly mitigated, because testing provided the 'ground truth' and the data necessary to calibrate the simulations through the intelligent use of tweaking 'knobs.'"

When testing was halted, the nation's Stockpile Stewardship Program came into being. Scientists now needed to prove that computer simulation results could hold their own and provide valuable information, which could be combined with data from current experiments and from underground tests to generate the necessary insights.

To bring such parity to computer simulations in the triumvirate of theory, experiment, and simulation, code designers had to address three concerns. First, could supercomputing hardware systems be built to perform the tasks? Second, could a workable simulation environment or support infrastructure be created for these systems? Third, could the mathematical algorithms used in the physics codes be scalable?

Bringing on the Hardware

The move to massively parallel processing supercomputers in the late 1980s was followed by the cessation of underground testing of nuclear devices in 1992 and the start of science-based stockpile stewardship.

The ASC Program required machines that could cost-effectively run simulations at trillions of operations per second (teraops) and use the terabytes of memory needed to properly express the complexity of the physics being simulated. This requirement forced a jump to massively parallel processing supercomputers that were, above all, scalable. In other words, these machines needed to be able to run large problems across the entire system without bogging down from communication bottlenecks, which led to the development of high-performance interconnects and the necessary software to manage these switches. Demands on hardware grew, and now the ASC Program at Livermore juggles three technology curves to ensure that users will have the machines they need today, tomorrow, and in the future. (See *S&TR*, June 2003, pp. 4–13.)

Creating an Infrastructure

Without a proper infrastructure, the ASC systems are little more than hard-to-program data-generation engines that create mind-numbing quantities of intractable, raw data. The infrastructure (sometimes called the supporting simulation environment) is what makes the terascale platform a real tool. The infrastructure includes improved systems software, input and output applications, message-passing libraries, storage systems, performance tools, debuggers, visualization clusters, data reduction and rendering algorithms, fiber infrastructure to offices, assessment theaters, high-resolution desktop displays, wide-area networks with encryption, and professional user consulting and services at the computer center—all focused on making the machines and codes run more efficiently.

The infrastructure has evolved in balance with the hardware. In 1999, for example, 2 terabytes of data from a three-dimensional simulation might have taken 2 or 3 days to move to archival storage or to a visualization server. By the end of 2000, that journey took 4 hours. Today, those 2 terabytes can zip from computer to mass storage in about 30 minutes. Similar efficiency and performance improvements have

director in the Computation Directorate. “From there, we ventured into scalar architecture and the massively parallel world of ASC supercomputers—systems of thousands of processors, each with a large supply of local memory. We were looking at not only sheer capability—which is the maximum processing power that can be applied to a single job—but also price performance. We were moving away from specialized processors for parallel machines to commodity processor systems and aggregating enough memory at

reasonable cost to address the new complexity and dimensionality.”

The challenge was to move into the world of massively parallel ASC systems in which thousands of processors may be working in concert on a problem. “First, we had to learn how to make these machines work at large scale,” says McCoy. “At the same time, we were developing massively parallel multiphysics codes and finding a way to implement them on the new machines. It was a huge effort in every direction.”

As the machines matured, the codes matured as well. “We’ve entered the young adult years,” says McCoy. “ASC White is running reliably in production mode, with a mean time to failure of a machine component measured in days, not hours or minutes. The proof-of-principle era is ending: The codes are deployed, the weapon designers increasingly are using these applications in major investigations, and this work is contributing directly to stockpile stewardship. With the upcoming 100-teraops ASC Purple, we believe that in many cases where we


occurred with compilers, debuggers, file systems, and data management tools as well as visualization and distance computing. Remote computing capabilities within the tri-laboratory community are easily available to all sites.

Designing Codes and Their Algorithms

Over the past few years, the ASC Program has developed some very capable three-dimensional codes and has maintained or further developed supporting science applications and two-dimensional weapons codes. Because of the enormous size of the computers and their prodigious power consumption, notes McCoy, the applications themselves are generally ignored by the media in favor of headline-producing computers. But if the truth were known, it is these codes and the people who build them, not the computers, that are the heart and soul of the ASC Program. “The computers come, and after a few years, they go,” says McCoy. “But the codes and code teams endure.” The greatest value of the ASC Program resides in these software assets, and this value is measured in billions of

dollars. The backbone of these scientific applications is mathematical equations representing the physics and the numerical constructs to represent the equations. To address issues such as how to handle a billion linear and nonlinear equations with a billion unknowns, computational mathematicians and others created innovative linear solvers (*S&TR*, December 2003, pp. 17–19) and Monte Carlo methods (*S&TR*, March 2004, pp. 19–21) that allow the mathematics to “scale” in a reasonable manner. Thus, as the problem grows more complex, processors can be added to keep the solution time manageable.

This composite image is taken from a three-dimensional simulation performed to help scientists better understand the sequence of events that led to the containment failure of the Baneberry underground test in December 1970. The complex geologic features of the Baneberry Event show pressure contours superimposed on geologic layers and faults. The different colors define different geologic material, including Paleozoic rock, tuff, and alluvium. The two turquoise lines slanting diagonally upward are faults that, scientists believe, contributed significantly to the containment failure. This simulation was the largest of its kind, requiring about 40 million zones, 3,000 time steps, and 40,000 processing hours on the Laboratory’s Multiprogrammatic Capability Resource machine, and created about 3 terabytes of data that had to be stored and interpreted.



have good experimental data, numerical error will be sufficiently reduced to make it possible to detect where physics models need improvement. We have demonstrated the value of high-resolution, three-dimensional physics simulations and are now integrating that capability into the Stockpile Stewardship Program, as we work to improve that capability by enhancing physics models. The ASC Program is no longer an initiative, it's a permanent element of a tightly integrated program with a critical and unambiguously defined national security mission."

Looking forward, Jim Rathkopf, an associate program leader for DNT's A Program, notes that with the arrival of Purple, codes will be able to use even higher resolution and better physics. "Higher resolution and better physics are required to reproduce the details of the different phases of a detonation and to determine the changes that occur in weapons as they age and their materials change over time."

Predicting Material Behavior

It's exciting times for scientists in the materials modeling world. The power of the terascale ASC machines and their codes

is beginning to allow physicists to predict material behavior from first principles—from knowing only the quantum mechanics of electrons and the forces between atoms. Earlier models, which were constrained by limited computing capabilities, had to rely on averages of material properties at a coarser scale than the actual physics demanded.

Elaine Chandler, who manages the ASC Materials and Physics Models Program, explains, "We can now predict very accurately the elastic properties of some metals. We're close to having predictive models for plastic properties as well."

Equation-of-state models are also moving from the descriptive to the predictive realm. It's possible to predict melt curves and phase boundaries from first principles and to predict changes in the arrangement of atoms from one crystalline structure to another. For example, scientists are running plasticity calculations to look at how tantalum moves and shears, then conducting experiments to see if their predictions are correct. Using this process, they can determine basic properties, such as yield strength.

With the older descriptive modeling codes, scientists would run many

experiments in differing regimes of temperature and pressure, then basically "connect the dots" to find out what a metal would do during an explosion. Now, they can perform the calculations that provide consistent information about the entire process. "It's a new world," says Chandler, "in which simulation results are trusted enough to take the place of physical experiments or, in some cases, lead to new experiments."

In the future, ASC Purple and the pioneering BlueGene/L computer will contribute to this new world. BlueGene/L is a computational-science research and evaluation machine that IBM will build in parallel with ASC Purple and deliver in 2005. According to Chandler, BlueGene/L should allow scientists to reach new levels of predictive capability for processes such as dislocation dynamics in metals, grain-scale chemical reactions in high explosives, and mixing in gases. Chandler says some types of hydrodynamics and materials science calculations will be relatively straightforward to port to the BlueGene/L architecture, but others, particularly those involving quantum-mechanical calculations,



(a) When the Cray-1 machine was installed in 1981, it was one of the fastest, most powerful scientific computers available. The last Cray obtained by Lawrence Livermore, in 1989, had 16 central processing units and about 2 megabytes of memory. (b) Nearly a decade later, the massively parallel 10-teraops ASC White arrived at the Laboratory as part of the National Nuclear Security Administration's Advanced Simulation and Computing Program. (For more on the history of supercomputing at Livermore, see *S&TR*, March 2002, pp. 20–26.)

Leaping from Milestone to Milestone

With the birth of the Stockpile Stewardship Program (SSP), the need for better computer simulations became paramount to help ensure that the nation's nuclear weapons stockpile remained safe, reliable, and capable of meeting performance requirements. The tri-laboratory (Livermore, Los Alamos, and Sandia national laboratories) Advanced Simulation and Computing (ASC) Program was created to provide the integrating simulation and modeling capabilities and technologies needed to combine new and old experimental data, past nuclear-test data, and past design and engineering experience. The first decade was devoted to demonstrating the proof of principle of ASC machines and codes. As part of that effort, the program set up a number of milestones to "prove out" the complex machines and their advanced three-dimensional physics codes.

The first milestone, accomplished in December 1999 by Livermore researchers on the ASC Blue Pacific/Sky machine, was the first-ever three-dimensional simulation of an explosion of a nuclear weapon's primary (the nuclear trigger of a hydrogen bomb). The simulation ran a total of 492 hours on 1,000 processors and used 640,000 megabytes of memory in producing 6 million megabytes of data contained in 50,000 computer files. The second Livermore milestone, a three-dimensional simulation of the secondary (thermonuclear) stage of a thermonuclear weapon, was accomplished in early 2001 on the ASC White machine—the first time that White was used to meet a milestone.

Livermore met a third milestone in late 2001, again using ASC White, coupling the primary and secondary in the first simulation of a full thermonuclear weapon. For this landmark simulation, the total run time was about 40 days of around-the-clock computing on over 1,000 processors. This simulation represented a major step toward deployment of the simulation capability. The quality was unusually high when compared to historic nuclear-test data. A detailed examination of the simulation results revealed complex coupled processes that had never been seen. In 2001, ASC White was also used by a Los Alamos team to complete an independent full-system milestone simulation.

In December 2002, Livermore completed another milestone on ASC White when a series of two-dimensional primary explosion calculations was performed. These simulations exercised new models intended to improve the physics fidelity and quantified the effect of increased spatial resolution on the accuracy of the results. The first production version of this code was also released at this time to users. Yet another Livermore team used ASC White to perform

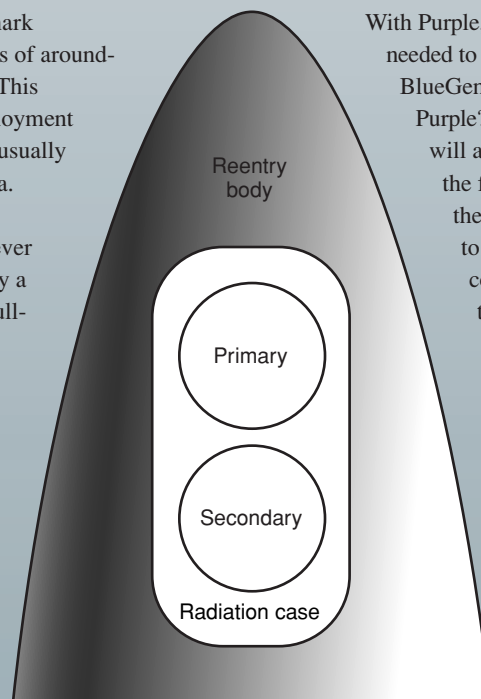
specialized three-dimensional simulations of a critical phase in the operation of a full thermonuclear weapon.

In 2003, Livermore teams completed separate safety and performance milestones. For the performance milestone, one team worked remotely on the ASC Q machine at Los Alamos to conduct a suite of three-dimensional primary explosion simulations in support of a Life Extension Program (LEP). Moving even farther from proof-of-principle demonstration and closer to deployment, a code team worked with the LEP team to accomplish this milestone, which addressed complex technical issues and contributed to meeting SSP objectives.

"We accomplished major objectives on time—with the early milestones demonstrating first-of-a-kind proof-of-principle capabilities," says Tom Adams, an associate program leader for DNT's A Program. "Achieving these milestones was the result of an intense effort by the code teams, who were assisted by dedicated teams from across the Laboratory. ASC milestones have now transitioned from these early demonstrations to milestones focused on improving the physics fidelity of the simulations and supporting stockpile stewardship activities. We are now in the position of delivering directly to the SSP."

Adams adds that the upcoming ASC Purple machine is a significant entry point. "Purple is the fulfillment of one of the original goals of the ASC Program, which is to bring a 100-teraops system to bear on stockpile stewardship issues. We need Purple to perform full, three-dimensional simulations for stockpile stewardship on a business-as-usual basis.

With Purple, we'll have the computing power and the codes needed to begin to address challenges in detail. Similarly, BlueGene/L will extend material models." And, beyond Purple? Petaops (quadrillion operations per second) systems will allow weapons designers and other users to address the fundamental underlying sources of uncertainty in the calculations. The goal is to be prepared to respond to technical issues that might arise because of component aging or new material requirements in the stockpile.



In thermonuclear weapons, radiation from a fission device (called a primary) can be contained and used to transfer energy for the compression and ignition of a physically separate component (called a secondary) containing thermonuclear fuel.


will require significant restructuring in order to use the architecture of this powerful machine. This is a challenge well worth the effort, because of the unprecedented computer power that BlueGene/L will offer to attack previously intractable problems.

“Nearly a half century ago,” adds Chandler, “scientists dreamed of a time when they could obtain a material’s properties from simply knowing the atomic numbers of the elements and quantum-mechanical principles. That dream eluded us because we lacked computers powerful enough to solve the complex calculations required. We are just now able to touch the edge of that dream, to reach the capabilities

needed to make accurate predictions about material properties.”

Delivering the Goods

ASC simulations play a key role in stockpile assessments and in programs to extend the life of the nation’s arsenal. Each year, a formal assessment reports the status of the nation’s stockpile of nuclear warheads and bombs. (See *S&TR*, July/August 2001, pp. 4–10.) This process involves the three national weapons laboratories working in concert to provide a “snapshot” of the stockpile’s health. Together, Livermore and Los Alamos are developing an improved methodology for quantifying confidence in the performance of these



A snapshot of dislocation microstructure generated in a massively parallel dislocation line dynamics simulation. The simulation was performed using the new code ParaDiS, developed at Livermore. ParaDiS has reached several computational milestones and outperformed all other existing codes by more than an order of magnitude in terms of computational complexity (that is, number of dislocation lines) and the magnitude of simulated plastic strain.

nuclear systems, with the goal of fully integrating this methodology into these annual assessments. The new methodology, known as quantification of margins and uncertainties (QMU), draws together information from simulations, experiments, and theory to quantify confidence factors for the key potential failure modes in every weapons system in the stockpile. (See *S&TR*, March 2004, pp. 19–21.)

The assertion that the nuclear explosive package in a weapon performs as specified is based on a design approach that provides an adequate margin against known potential failure modes. Weapons experts judge the adequacy of these margins using data from past nuclear experiments, ground and flight tests, and material compatibility evaluations during weapons development as well as routine stockpile surveillance, nonnuclear tests, and computer simulations. With the cessation of underground nuclear testing, the assessment of these margins relies much more heavily on surveillance and computer simulations than in the past and therefore requires the simulations to be more rigorous and detailed.

Because no new weapons are being developed, the existing ones must be maintained beyond their originally planned lifetimes. To ensure the performance of these aging weapons, Livermore and Los Alamos weapons scientists use QMU to help them identify where and when they must refurbish a weapons system. When needed, a Life Extension Program (LEP) is initiated to address potential performance issues and extend the design lifetime of a weapons system through refurbishment or replacement of parts. For the W80 LEP now under way, results from ASC simulations are weighed along with data from past nuclear weapons tests and from recent small-scale science tests. These results will support certification of the LEP.

Using today’s ASC computer systems and codes, scientists can include unprecedented geometric fidelity in addressing issues specific to life extension. They can also investigate particular aspects, such as

plutonium's equation of state, scientifically and in detail, and then extend that understanding to the full weapons system. The results of these simulations, along with data from legacy testing and current experiments, improve the ability of weapons designers to make sound decisions in the absence of nuclear testing.

As computational capability increases, designers will have a more detailed picture of integrated weapons systems and can address even more complex issues—for example, how various materials fracture—with even higher resolution.

Right Answers for Right Reasons

Even as inaccuracies due to mathematics and numerics are being resolved by running simulations at ever higher resolutions, the question remains: If a simulation result is unusual, how do scientists know whether it is a problem due to inadequate resolution or simply an error, or bug, in the coding?

According to Cynthia Nitta, manager of the ASC Validation and Verification (V&V) Program, an effort was established by the ASC Program to rigorously examine the computational science and engineering simulation results with an eye to their credibility. "Can we trust that the results of simulations are accurate? Do the results reflect the real-world phenomena that they are striving to re-create or predict?" asks Nitta. "In the V&V Program, we are developing a process that should increase the confidence level for decisions regarding the nation's nuclear stockpile. Our methods and processes will establish that the calculations provide the right answers for the right reasons."

The verification process determines whether a computer simulation code for a particular problem accurately represents the solutions of the mathematical model. Evidence is collected to ascertain whether the numerical model is being solved correctly. This process ensures that sound software-quality practices are used and the software codes themselves are free of defects and errors. It also checks that the

code is correctly solving the mathematical equations in the algorithms and verifies that the time and space steps or zones chosen for the mathematical model are sufficiently resolved.

The validation process determines whether the mathematical model being used accurately represents the phenomenon being modeled and to what degree of accuracy. This process ensures that the simulation adequately represents the appropriate physics by comparing the output of a simulation with data gathered in experiments and quantifying the uncertainties in both.


Nitta says, "Computer simulations are used in analyzing all aspects of weapons

systems as well as for analyzing and interpreting weapons-related experiments. The credibility of our simulation capabilities is central to the credibility of the certification of the nuclear stockpile. That credibility is established through V&V analyses."


Terascale—A Beginning, Not an End

With the proof-of-principle phase ending and new codes being deployed, what does the future hold? With the arrival of the 100-teraops Purple in 2005, many simulations become possible, including a full-system calculation of a nuclear weapon with sufficient resolution to distinguish between phenomenological and numerical issues. But, as McCoy, Christensen, and


(a) Simulation



(b) Experimentation



(c) Quantitative comparison



The verification and validation (V&V) process ties together (a) simulations and (b) experiments using (c) quantitative comparisons.



The Terascale Simulation Facility (TSF) at Lawrence Livermore will have two machine rooms for housing ASC Purple and BlueGene/L. The 100-teraops Purple will be powered by 12,288 microprocessors (in 1,536 individual nodes) and have 50 terabytes of memory (400,000 times more capacity than the average desktop computer) and 2 petabytes of disk storage (the content of approximately one billion books). The facility also features substantial visualization areas, including theaters and laboratories to handle the extremely large data sets produced.

others point out, 100 teraops is just the beginning.

The ASC Program plans to increase the level of confidence in predictions that such simulations can bring as well as increase the predictive capability, by tying together simulations and experiments even more closely and quantifying the uncertainty of the simulated results.

“We’re positioning our science codes to run on the Purple and BlueGene/L machines so that we can understand the physics in even greater detail,” says

Christensen. “It’s been a challenging journey over the past decade: In the ASC Program, we’ve demonstrated that we can acquire and use the world’s most powerful computers to perform three-dimensional calculations that capture many details of weapons performance. Now, we must look toward the next goal, which is to be able to predict weapons behavior and quantify the confidence we have in that prediction. If the past decade is any indication—and we believe it is—this is a goal we can, and will, indeed attain.”

—Ann Parker

Key Words: Accelerated Strategic Computing Initiative (ASCI), Advanced Simulation and Computing (ASC) Program, ASC Purple, ASC White, BlueGene/L, descriptive models, Life Extension Program, materials modeling, predictive models, quantification of margins and uncertainties (QMU), Stockpile Stewardship Program (SSP), Verification and Validation (V&V) Program, weapons certification.

For further information contact Randy Christensen (925) 423-3054 (christensen5@llnl.gov).